

Monotone projection lower bounds from extended formulation lower bounds

Joshua A. Grochow*

November 4, 2015

Abstract

In this short note, we show that the Hamilton Cycle polynomial, $\sum_{n\text{-cycles } \sigma} \prod_{i=1}^n x_{i,\sigma(i)}$ is not a monotone sub-exponential-size projection of the permanent; this both rules out a natural attempt at a monotone lower bound on the Boolean permanent, and shows that the permanent is *not* complete for non-negative polynomials in $\text{VNP}_{\mathbb{R}}$ under monotone p-projections. We also show that the cut polynomials, $\sum_{A \subseteq [n]} \prod_{i \in A, j \notin A} x_{ij}^q$, and the perfect matching polynomial (or “unsigned Pfaffian”) $\frac{1}{2^n n!} \sum_{\pi \in S_{2n}} \prod_{i=1}^n x_{\pi(2i-1), \pi(2i)}$ are not monotone p-projections of the permanent. The latter can be interpreted as saying that there is no monotone projection reduction from counting perfect matchings in general graphs to counting perfect matchings in bipartite graphs, putting at least one theorem behind the well-established intuition. To prove these results we introduce a new connection between monotone projections of polynomials and extended formulations of linear programs that may have further applications.

1 Introduction

The permanent $\text{perm}_n(X) = \sum_{\pi \in S_n} x_{1,\pi(1)} x_{2,\pi(2)} \cdots x_{n,\pi(n)}$ has long-fascinated combinatorialists [Min84, vLW01, MM60], more recently physicists [WS10, AA11], and since Valiant’s seminal paper [Val79b], has also been a key object of study in computational complexity. Despite its beauty, the permanent has some computational quirks: in particular, although the permanent of integer matrices is $\#\text{P}$ -complete and the permanent is VNP -complete in characteristic zero, the permanent *mod* 2 is the same as the determinant, and hence can easily be computed. In fact, computing the permanent *mod* 2^k is easy for any k [Val79b], though the proof is more involved. Modulo any other number n , the permanent of integer matrices is $\text{Mod}_n \text{P}$ -complete.

In contrast, the seemingly similar Hamilton Cycle polynomial,

$$HC_n(X) = \sum_{n\text{-cycles } \sigma} x_{1,\sigma(1)} x_{2,\sigma(2)} \cdots x_{n,\sigma(n)},$$

where the sum is only over n -cycles rather than over all permutations, does not have these quirks: The Hamilton Cycle polynomial is VNP -complete over any ring R [Val79a] and $\text{Mod}_n \text{P}$ -complete for all n (that is, counting Hamilton cycles is complete for these Boolean counting classes).

*Santa Fe Institute, Santa Fe, NM, USA, jgrochow@santafe.edu

Jukna [Juk14b] observed that, over the Boolean semiring, if the Hamilton Cycle polynomial were a monotone p-projection of the permanent, there would be a $2^{n^{\Omega(1)}}$ lower bound on monotone circuits computing the permanent, a lower bound that still remains open (the current record is still Razborov’s $n^{\Omega(\log n)}$ [Raz85]). Even over the real numbers, such a monotone p-projection would give an alternative proof of a $2^{n^{\Omega(1)}}$ lower bound on the permanent (Jerrum and Snir [JS82] already showed the permanent requires monotone circuits of size $2^{\Omega(n)}$ over the \mathbb{R} and over the tropical semiring). Here we show that no such monotone reduction exists—over \mathbb{R} nor over the Boolean semiring—by connecting monotone p-projections to extended formulations of linear programs.

We use the same technique to show that the perfect matching polynomial or “unsigned Pfaffian”

$$\frac{1}{2^n n!} \sum_{\pi \in S_{2n}} \prod_{i=1}^n x_{\pi(2i-1), \pi(2i)} = \sum_{\substack{\pi \in S_{2n} \\ \pi(1) < \pi(3) < \dots < \pi(2n-1) \\ \pi(2k-1) < \pi(2k) \ \forall k}} \prod_{i=1}^n x_{\pi(2i-1), \pi(2i)}$$

is not a monotone p-projection of the permanent. As the perfect matching polynomial counts perfect matchings in a general graph, and the permanent counts perfect matchings in a bipartite graph, we have:

Corollary 1. *Any efficient parsimonious projection reduction from counting perfect matchings in general graphs to counting perfect matchings in bipartite graphs must be non-monotone (and, therefore, seemingly quite unintuitive!).*

Remark 1 (On the Boolean semi-ring). Our results also hold for *formal polynomials* over the Boolean semi-ring $\mathbb{B} = (\{0, 1\}, \vee, \wedge)$. Over the Boolean semi-ring, the permanent is the indicator function of the existence of a perfect matching in a bipartite graph, and the unsigned Pfaffian is the indicator function of a perfect matching in a general graph. However, over \mathbb{B} , each function is represented by more than one formal polynomial, and we do not yet know how to extend our results to the setting of *functions* over \mathbb{B} . See Section 5 for details and specific questions.

We also use the same technique to show that the cut polynomials $\text{Cut}^q = \sum_{A \subseteq [n]} \prod_{i \in A, j \notin A} x_{ij}^q$ are not monotone p-projections of the permanent. Perhaps the main complexity-theoretic interest in the cut polynomials is that Cut^q over the finite field \mathbb{F}_q is the only known example of a natural polynomial that is neither in $\text{VP}_{\mathbb{F}_q}$ nor $\text{VNP}_{\mathbb{F}_q}$ -complete under a standard complexity-theoretic assumption (that PH doesn’t collapse) [Bür99]; there it was also shown that if $\text{VP}_{\mathbb{F}_q} \neq \text{VNP}_{\mathbb{F}_q}$ then such polynomials of intermediate complexity must exist. In that paper, it was asked whether the cut polynomials, considered as polynomials over the rationals, were $\text{VNP}_{\mathbb{Q}}$ -complete. Although our results don’t touch on this question, these previous results motivate the study of these polynomials over \mathbb{Q} .

Finally, we note that our results shed a little more light on the complicatedness of the known VNP -completeness proofs for the permanent [Val79b, Aar11]. Namely, prior to our result, the fact that the permanent is not hard modulo 2 already implied that any completeness result must use 2 in a “bad” way: for example, dividing by 2 somewhere, or somewhere requiring that some quantity that is necessarily even be non-zero. This is indeed true both of Valiant’s original proof [Val79b] and of Aaronson’s independent quantum linear-optics proof [Aar11]. One might hope for a classical analogue of Aaronson’s quantum proof, using the characterization of BPP in terms of stochastic matrices as a replacement

for the characterization of BQP using unitary matrices. However, our result says that any completeness proof for the permanent must use non-monotone reductions, so such a classical analogue is not possible:

Corollary 2. *Aaronson’s quantum linear optics proof [Aar11] that the permanent is #P-hard cannot be replaced by one using classical randomized algorithms in place of quantum algorithms.*

In light of these results, Valiant’s 4×4 gadget may perhaps seem less mysterious than the fact that such a gadget exists that is *only* 4×4 !

To prove these results, we show that a monotone projection between non-negative polynomials essentially implies that the Newton polytope of one polynomial is an extension of the Newton polytope of the other (Lemma 1), and then apply known lower bounds on the extension complexity of certain polytopes. We hope that the connection between Newton polytopes, monotone projections, and extended formulations finds further use.

2 Preliminaries

A polynomial $f(x_1, \dots, x_n)$ is a (*simple*) *projection* of a polynomial $g(y_1, \dots, y_m)$ if f can be constructed from g by replacing each y_i with a constant or with some x_j . The polynomial f is an *affine projection* of g if f can be constructed from g by replacing each y_i with an affine linear function $\ell_i(\vec{x})$. When we say “projection” we mean simple projection. Given two families of polynomials (f_n) , (g_n) , if there is a function $p(n)$ such that f_n is a projection of $g_{p(n)}$ for all sufficiently large n , then we say that (f_n) is a projection of (g_n) with *blow-up* $p(n)$. If (f_n) is a projection of (g_n) with polynomial blow-up, we say that (f_n) is a *p-projection* of (g_n) .

Over any subring of \mathbb{R} —or more generally any totally ordered semi-ring—a *monotone projection* is a projection in which all constants appearing in the projection are non-negative.¹ Monotone p-projection is defined analogously.

To each monomial $x_1^{e_1} \cdots x_n^{e_n}$ we associate its exponent vector (e_1, \dots, e_n) , as a point in $\mathbb{N}^n \subseteq \mathbb{R}^n$. We then have:

Definition. The *Newton polytope* of a polynomial $f(x_1, \dots, x_n)$, denoted $\text{New}(f)$, is the convex hull in \mathbb{R}^n of the exponent vectors of all monomials appearing in f with non-zero coefficient.

A polytope is *integral* if all its vertices have integer coordinates; note that Newton polytopes are always integral.

For a polytope P , let $c(P)$ denote the “complexity” of P , as measured by the minimal number of linear inequalities needed to define P . A polytope $Q \subseteq \mathbb{R}^m$ is an extension of $P \subseteq \mathbb{R}^n$ if there is an affine linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that $\pi(Q) = P$. The *extension complexity* of P , denoted $xc(P)$, is the minimum complexity of any extension of P (of any dimension): $xc(P) = \min\{c(Q) \mid Q \text{ is an extension of } P\}$.

¹Over the Boolean semi-ring, we must restrict to projections that only use variables and not their negations. This is because, in some contexts, general Boolean projections may allow literals and not just variables, since otherwise there is no way to achieve negation in a projection using only \vee and \wedge .

3 Main Lemma

Lemma 1. *Let R be a totally ordered semi-ring, and let $f(x_1, \dots, x_n)$ and $g(y_1, \dots, y_m)$ be polynomials over R with non-negative coefficients. If f is a monotone projection of g , then the intersection of $\text{New}(g)$ with some linear subspace is an extension of $\text{New}(f)$. In particular, $xc(\text{New}(f)) \leq m + c(\text{New}(g))$.*

Recall that a *term* of a polynomial is a monomial together with its coefficient.

Proof. Under simple projections, a monomial in the y 's maps to some scalar multiple of a monomial in the x 's (possibly the empty monomial, resulting in a constant term, or possibly the zero multiple, resulting in zero). Let π be a monotone projection map, defined on the variables y_i , and extended naturally to monomials and terms in the y 's. Since each term t of g is a monomial multiplied by a positive coefficient, and since π is non-negative, $\pi(t)$ is either zero or a single monomial in the x 's with nonzero coefficient. The former situation can happen only if t contains some variable y_i such that $\pi(y_i) = 0$. Let $\ker(\pi)$ denote the set $\{y_i | \pi(y_i) = 0\}$. Thus, for every term t of g that is disjoint from $\ker(\pi)$, $\pi(t)$ actually appears (possibly with a different coefficient, but still non-zero) in f , since no two terms can cancel under projection by π .

Let e_1, \dots, e_m be the coordinates on \mathbb{R}^m , the ambient space of $\text{New}(g)$. Let K denote the subspace of \mathbb{R}^m defined by the equations $e_i = 0$ for each i such that $y_i \in \ker(\pi)$. Let P be the intersection of $\text{New}(g)$ with K , considered as a polytope in K ; note that P is exactly the convex hull of the exponent vectors of monomials in g that are disjoint from $\ker(\pi)$. Since π is multiplicative on monomials, it induces a *linear* map ℓ_π from K to \mathbb{R}^n (the ambient space of $\text{New}(f)$). By the previous paragraph, the exponent vectors of f are exactly ℓ_π applied to the exponent vectors of monomials in g that are disjoint from $\ker(\pi)$. By the linearity of ℓ_π and the convexity of P and $\text{New}(f)$, we have that $\text{New}(f) = \ell_\pi(P)$, so P is an extension of $\text{New}(f)$. Since P is defined by intersecting $\text{New}(g)$ with $\leq m$ additional linear equations, the lemma follows. \square

Several partial converses to our Main Lemma also hold. Perhaps the most natural and interesting of these is:

Observation 1. *Given any sequence of integral polytopes $(P_n \subseteq \mathbb{R}^n)$ such that the poly(n)-th cycle cover polytope is an extension of P_n along a projection $\pi_n: \mathbb{R}^{\text{poly}(n)} \rightarrow \mathbb{R}^n$ with integer coefficients of polynomial bit-length, there is a sequence of polynomials $(f_n) \in \text{VNP}$ such that $\text{New}(f_n) = P_n$ and f is a monotone p -projection of the permanent.*

This construction works over any semi-ring of “characteristic zero,” that is, in which $1 + \dots + 1 \neq 0$ (any number of times); this includes the Boolean semi-ring, since there the addition operation is \vee (not \oplus).

Proof. Let C_m denote the m -th cycle cover polytope, let $m(n)$ be a polynomial such that $C_{m(n)}$ is an extended formulation of P_n , and let $b(n)$ be a polynomial upper bound on the bit-length of the coefficients of π_n . Let V_m denote the vertex set of the cycle cover polytope, i.e. the incidence vectors of cycle covers. Define f_n as $\sum_{\vec{e} \in V_m} \vec{y}^{\pi(\vec{e})}$, where $\vec{y} = (y_1, \dots, y_n)$. As every exponent vector of f_n is in $\pi_n(C_{m(n)}) = P_n$, and conversely every vertex of P_n is an exponent vector of f_n ,² we have $\text{New}(f_n) = P_n$. Furthermore, f_n is a monotone *nonlinear* projection of the permanent using the map $x_{ij} \mapsto \vec{y}^{\pi((0,0,\dots,1,\dots,0))}$, where the 1 is in the (i, j)

²This is the only place we need the assumption of characteristic zero; see the remark before the proof.

position. Using the universality of the permanent and repeated squaring, this can easily be turned into a monotone *simple* projection of the permanent of size $\text{poly}(m(n), b(n))$. \square

This can be generalized from the cycle cover polytopes and the permanent to arbitrary integral polytopes and the natural associated polynomial (the sum over all monomials whose exponent vectors are vertices of the polytope), but at the price of using “monomial projections”—in which each variable is replaced by a monomial—rather than simple projections. There ought to be a version of this observation over sufficiently large fields and allowing rational coefficients in π and using Strassen’s division trick [Str73], but the only such versions the author could come up with had so many hypotheses as to seem uninteresting.

4 Applications

Note: The following theorems hold over any totally ordered semi-ring, including the Boolean and-or semi-ring, the non-negative real numbers under addition and multiplication, and the tropical semi-ring of real numbers under min and addition.

Theorem 3. *The Hamilton Cycle polynomial is not a monotone affine p -projection of the permanent; in fact, any monotone affine projection from the permanent to the Hamilton Cycle polynomial has blow-up at least $2^{\Omega(n)}$.*

Proof. First, recall that if an n -variable polynomial is an affine projection of the $m \times m$ permanent, then it is a simple projection of the $(n+1)m \times (n+1)m$ permanent. For completeness we recall the brief proof: Let $\ell_{ij}(\vec{x})$ be the affine linear function corresponding to the variable y_{ij} of the $m \times m$ permanent, and write $\ell_{ij} = a_0 + a_1x_1 + \dots + a_nx_n$. Let G be the complete directed graph with loops on m vertices and edge weights y_{ij} . Replace the edge (i, j) by $n+1$ parallel edges with weights $a_0, a_1x_1, \dots, a_nx_n$. Add a new vertex on each of these parallel edges, splitting each parallel edge into two. For the edge weighted a_0 , the two edges have weights $1, a_0$, and for the remaining edges the new edges get weights a_i, x_i . It is a simple and instructive exercise to see that this has the desired effect.

Now we show the result for simple projections. If the Hamilton Cycle polynomial were a monotone projection of the permanent, then by the Main Lemma, $\text{New}(\text{perm})$ (intersected with a linear subspace) would be an extension of $\text{New}(HC)$.

The Newton polytope of the permanent is the convex hull of all vectors in $\{0, 1\}^{n^2}$ corresponding to directed cycle covers of a graph, as each monomial in the permanent corresponds to such a cycle cover. The cycle cover polytope can easily be described by the n equations saying that each vertex has in-degree and out-degree exactly 1. Thus $c(\text{New}(\text{perm}_n)) \leq n$.

But the Newton polytope of the Hamilton Cycle polynomial is exactly the TSP polytope, which by [Rot14, Corollary 2] requires extension complexity $2^{\Omega(n)}$. \square

Theorem 4. *The perfect matching polynomial (or “unsigned Pfaffian”) is not a monotone affine p -projection of the permanent; in fact, any monotone affine projection from the permanent to the perfect matching polynomial has blow-up at least $2^{\Omega(n)}$.*

Proof. The proof is the same as for the Hamilton Cycle polynomial, using [Rot14, Theorem 1], which gives a lower bound of $2^{\Omega(n)}$ on the extension complexity of the perfect matching polytope, which is the Newton polytope of the perfect matching polynomial. \square

Theorem 5. *For any q , the q -th cut polynomial is not a monotone affine p -projection of the permanent; in fact, any monotone affine projection from the permanent to the q -th cut polynomial has blow-up at least $2^{\Omega(n)}$.*

Proof. Use [FMP⁺12, Theorem 7], which says that $xc(\text{New}(\text{Cut}^1)) \geq 2^{\Omega(n)}$, as $\text{New}(\text{Cut}^1)$ is the cut polytope. The one additional observation we need is that $\text{New}(\text{Cut}^q)$ is just the q -scaled version of $\text{New}(\text{Cut}^1)$, and this rescaling does not affect the extension complexity. \square

5 Open Questions

Despite the common feeling that Razborov’s super-polynomial lower bound [Raz85] on monotone circuits for CLIQUE “finished off” monotone Boolean circuit lower bounds, several natural and interesting question remain. For example, does Directed s - t Connectivity require monotone Boolean circuits of size $\Omega(n^3)$? (A matching upper bound is given by the Bellman–Ford algorithm.) Is there a monotone Boolean reduction from general perfect matching to bipartite perfect matching? A positive answer to the following question would rule out such monotone (projection) reductions:

Open Question 1. *Extend Theorem 4 from formal polynomials over the Boolean semi-ring to Boolean functions.*

However, there are even easier questions, intermediate between the Boolean function case and the algebraic case considered in this paper; Jukna [Juk14a] discusses the notion of one polynomial “counting” another, which means that they agree on all $\{0, 1\}$ inputs.

Open Question 2. *Prove that no monotone polynomial-size projection of the permanent agrees with the perfect matching polynomial on all $\{0, 1\}$ inputs (“counts the perfect matching polynomial”). Similarly, prove that no monotone polynomial-size projection of the permanent counts the Hamilton cycle polynomial.*

S. Jukna asked (personal communication) whether any monotone polynomial-size projection of the s - t connectivity polynomial counts the Hamilton path polynomial. A positive answer would give a direct proof that no monotone polynomial-size circuit counts the s - t connectivity polynomial; the only current proof we are aware of [Juk14a] goes through Razborov’s lower bound on CLIQUE [Raz85], followed by Valiant’s reduction from the clique polynomial to the Hamilton path polynomial [Val79a], followed by a standard reduction from Hamilton path to counting s - t paths. One natural approach to a positive answer would be to show that the Hamilton path polynomial is in fact a monotone p -projection of the s - t connectivity polynomial. The following question could rule out such an approach, which would then suggest why such a roundabout proof was needed for the lower bound on counting s - t connectivity:

Open Question 3 (S. Jukna, personal communication). *Is the m -th s - t path polytope an extension of the n -th TSP polytope with $m \leq \text{poly}(n)$?*

Since the separation problem for the s - t path polytope is NP-hard (see, e.g., [Sch03, §13.1]), answering this question negatively seems to require more subtle understanding of these polytopes than “simply” an extended formulation lower bound.

Another example of a natural polytope question with a similar flavor comes from the cut polynomials. In combination with Bürgisser’s results and questions on the cut polynomials [Bür99] (discussed in Section 1), we are led to the following question:

Open Question 4. *Is the m -th cut polytope an extension of the n -th TSP polytope, for $m \leq \text{poly}(n)$?*

A negative answer would show that Cut^q is not complete for non-negative polynomials in $\text{VNP}_{\mathbb{Q}}$ under monotone p -projections, though as with the example of the permanent, this is not necessarily an obstacle to being VNP -complete under general p -projections. Yet even the monotone completeness of the cut polynomials remains open. In fact, even more basic questions remain open:

Open Question 5. *Is every non-negative polynomial in VNP a monotone projection of the Hamilton Cycle polynomial? Is there any polynomial that is “positive VNP -complete” in this sense?*

To relate this to the current proofs of VNP -completeness of HC_n , we need to draw a distinction. Let $\text{VP}_{\mathbb{R}}^{\geq 0}$ denote the polynomial families in $\text{VP}_{\mathbb{R}}$ all of whose coefficients are non-negative, and let $\text{mVP}_{\mathbb{R}}$ (“monotone VP ”) denote the class of families of polynomials with polynomially many variables, of polynomial degree, and computable by polynomial-size *monotone* circuits over \mathbb{R} . Define $\text{VNP}_{\mathbb{R}}^{\geq 0}$ to be the non-negative polynomials in $\text{VNP}_{\mathbb{R}}$, and $\text{mVNP}_{\mathbb{R}}$ to be the function families of the form $f_n = \sum_{\vec{e} \in \{0,1\}^{\text{poly}(n)}} g_m(\vec{e}, \vec{x})$, where $m \leq \text{poly}(n)$ and $(g_m) \in \text{mVP}_{\mathbb{R}}$.

Valiant’s original completeness proof for the Hamilton Cycle polynomial [Val79a] is “mostly” monotone: It uses polynomial-size formulas for the coefficients of the monomials (coming from the definition of VNP), but otherwise is entirely monotone. In other words, the proof shows that HC is mVNP -hard under monotone projections. However, we note that it’s not clear whether HC is even *in* mVNP ! Question 5 asks whether HC , or indeed any polynomial, is $\text{VNP}^{\geq 0}$ -complete under monotone projections; the question of whether there exist polynomials that are mVNP -complete under monotone projections also seems potentially interesting.

Finally, we ask about stronger notions of monotone reduction, which seem to require a different kind of proof technique. Recall that a *c-reduction* from f to g is a family of algebraic circuits for f with oracle gates for g .

Open Question 6. *Do the analogues of Theorems 3–5 hold for monotone weakly-skew c -reductions in place of affine p -projections? What about monotone general c -reductions?*

Acknowledgment. We would like to thank Stasys Jukna for the question that motivated this paper [Juk14b], and `cstheory.stackexchange.com` for providing a forum for the question. We also thank Stasys for comments on a draft and interesting discussions leading to Questions 2 and 3. We thank Leslie Valiant for an interesting conversation that led to Question 5. We thank Ketan Mulmuley and Youming Qiao for collaborating on [GMQ15], which is why the author had Newton polytopes on the mind.

References

- [AA11] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC ’11, pages 333–342, New York, NY, USA, 2011. ACM.
- [Aar11] Scott Aaronson. A linear-optical proof that the permanent is $\#P$ -hard. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 467(2136):3393–3405, 2011.

- [Bür99] Peter Bürgisser. On the structure of Valiant’s complexity classes. *Discrete Math. Theor. Comput. Sci.*, 3(3):73–94, 1999.
- [FMP⁺12] Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pages 95–106, New York, NY, USA, 2012. ACM.
- [GMQ15] Joshua A. Grochow, Ketan D. Mulmuley, and Youming Qiao. Degenerations of VP and VNP. In preparation, 2015.
- [JS82] Mark Jerrum and Marc Snir. Some exact complexity results for straight-line computations over semirings. *J. Assoc. Comput. Mach.*, 29(3):874–897, 1982.
- [Juk14a] Stasys Jukna. Lower bounds for monotone counting circuits. ECCC Tech. Report TR14-169, 2014.
- [Juk14b] Stasys Jukna. “Why is Hamilton Cycle so different from permanent?”. <http://cstheory.stackexchange.com/q/27496/129>, 2014.
- [Min84] Henryk Minc. *Permanents*. Cambridge University Press, 1984.
- [MM60] Thomas Muir and William H. Metzler. *A treatise on the theory of determinants*. Dover, New York, 1960.
- [Raz85] A. A. Razborov. Lower bounds of monotone complexity of the logical permanent function. *Mat. Zametki*, 37(6):887–900, 942, 1985.
- [Rot14] Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC ’14, pages 263–272, New York, NY, USA, 2014. ACM. Preprint available as arXiv:1311.2369 [cs.CC].
- [Sch03] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Paths, flows, matchings, Chapters 1–38.
- [Str73] Volker Strassen. Vermeidung von Divisionen. *J. Reine Angew. Math.*, 264:184–202, 1973.
- [Val79a] Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC ’79, pages 249–261, New York, NY, USA, 1979. ACM.
- [Val79b] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.
- [vLW01] J. H. van Lint and R. M. Wilson. *A course in combinatorics*. Cambridge University Press, 2001.
- [WS10] Tzu-Chieh Wei and Simone Severini. Matrix permanent and quantum entanglement of permutation invariant states. *Journal of Mathematical Physics*, 51(9), 2010.